DOCUMENT RESUME

ED 341 026 CS 010 801

AUTHOR Kantrov, Ilene
TITLE Talking to the Computer: A Prototype Speech Recognition System for Early Reading Instruction. Report No. 91-3.
INSTITUTION Education Development Center, Inc., Newton, Mass.
PUB DATE 91
NOTE 14p.
PUB TYPE Reports - Descriptive (141) -- Reports - Research/Technical (143)

EDRS PRICE MF01/PC01 Plus Postage.
DESCRIPTORS *Beginning Reading; *Computer Assisted Instruction; *Computer Software; Grade 1; Kindergarten; Kindergarten Children; Man Machine Systems; Primary Education; Reading Instruction; Young Children
IDENTIFIERS Prototypes; *Speech Recognition

ABSTRACT

While computer software that talks has become commonplace, software that listens still remains quite rare. A demonstration speech recognition program designed for early reading instruction developed in the mid-1980s never reached the market due to the immaturity of the technology and the volatility of the microcomputer hardware and software business. The program used an interactive speech technology system (providing meaningful and appealing contexts in which young children explored an 18-word vocabulary list) that did not make many errors in recognition, compared to other low-cost systems. The four components of the software were: a training program; a story program; three games designed to help children learn vocabulary; and a program that allowed parents or teachers to enter additional vocabulary words. Three field tests of the system over a period of 6 months were conducted with 42 kindergarten and first-grade children in 2 Boston-area public schools. Results indicated that: (1) despite technical problems, children enjoyed using the program; (2) the opportunity for children to control what happened on the screen by reading words aloud was a powerful motivation to learn to read; and (3) the program was most successful with children who were just beginning to recognize words and with those having difficulty learning to read. As the prices of speech recognition systems come down and quality goes up, the experiences of this demonstration speech recognition program may prove relevant to others beginning to explore speech recognition's potential for reading instruction and remediation. (Two figures representing video images produced by the program and three tables of data are included.) (RS)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

REPORTS AND PAPERS IN PROGRESS

ED341026

SCOPE OF INTEREST NOTICE
The ERIC Facility has assigned this document for processing to: CS
In our judgment, this document is also of interest to the Clearinghouses noted to the right. Indexing should reflect their special points of view. SE

Talking to the Computer

A Protptype Speech Recognition System
for Early Reading Instruction

Ilene Kantrov

Report No. 91-3

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY
Leigh Lake
TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)
☑ This document has been reproduced as received from the person or organization originating it.
☐ Minor changes have been made to improve reproduction quality.
• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy

CS010801

EDC
EDUCATION DEVELOPMENT CENTER
55 CHAPEL STREET • NEWTON, MA • 02160

BEST COPY AVAILABLE

Talking to the Computer

A Prototype Speech Recognition System for Early Reading Instruction

Ilene Kantrov

Report No. 91-3

Publications in the series Reports and Papers in Progress grow out of the work of the over twenty projects of the Center for Learning, Teaching, and Technology (LTT) of Education Development Center. LTT was established in 1982 as the Center for Learning Technology to study and develop applications for information technologies. Early work explored new and productive uses in education for microcomputers, video, CD-ROM, and multimedia. Building on this work with technology, a talented and skilled staff of fifty currently engage in research and policy, product design and development, professional development programs, and dissemination. LTT now represents a wide-ranging capacity that spans all levels from early childhood to college and in-service, a variety of content domains, and learning and teaching in and out of schools.

The Center for Learning, Teaching, and Technology is part of Education Development Center, Inc. (EDC), an international, nonprofit, research and development organization with offices in Newton, Massachusetts, and Washington, D.C. For over three decades, EDC has been recognized as a leading center for curriculum reform, training, research, and institutional development. EDC works throughout the United States and around the world to solve a wide range of educational, health, and social problems.

To receive a list of current publications, please write to

Publications
Center for Learning, Teaching, and Technology
Education Development Center
55 Chapel Street
Newton, MA 02160

Talking to the Computer

A Prototype Speech Recognition System for Early Reading Instruction

Ilene Kantrov

This paper reports on the field tests of one speech recognition system for early reading instruction. The software expanded the use of speech recognition technology beyond simple drill and practice by providing meaningful and appealing contexts in which young children (here, kindergarteners and first-graders) explored an eighteen-word vocabulary list. The results of these field tests reveal both the particular challenges involved in designing educational programs using speech recognition technology, and promising directions for future work.

Computer software that talks—through speech synthesis or digitized speech—has become commonplace (Parham 1988). We now have a variety of talking word processors and programs that issue spoken instructions, as well as programs that read stories and speak the dialogue of plays.

Computer software that *listens*, on the other hand, still remains quite rare. The technological and economic barriers to good quality speech recognition at an affordable price remain formidable. A number of products have appeared on the market, but none has met with great commercial success, and nearly all have been aimed at the business rather than the school market .[1] The few speech recognition products targeted at schools,[2] such as those produced by Chatterbox Voice Learning Systems, focus on drill in reading and math. The Chatterbox Voice Reading Ability

Drill software, for example, drills students on lists of words (from a total vocabulary of 800 words), providing feedback for correct and incorrect responses.[3]

Yet the educational possibilities of speech recognition technology go beyond drill and practice. Hopes and expectations for significant applications of speech recognition have been especially high for reading instruction (Blanchard et al. 1987; Strickland et al. 1987). This report describes a set of software programs for early reading instruction developed in the mid-1980s by me and my colleagues at Education Development Center in Newton, Massachusetts. The software was designed to demonstrate some of the capabilities—beyond simple drill and practice—of speech recognition, combined with speech output, for young children learning to read. We were particularly interested in using the technology to help children learn to read words in meaningful and appealing contexts, as many reading specialists advocate.

Building on the original set of demonstration programs, we had hoped to develop a series of speech-based reading software that would take full advantage of the possibilities of the technology. Because of the immaturity of the technology, however, as well as the volatility in the microcomputer hardware and software business at the time (1984-85), our demonstration programs never reached the market and we never got to develop additional software. In the 1990s, as prices of speech recognition systems come down and quality goes up, our experience with the technology may prove relevant to others beginning to explore speech recognition's potential for reading instruction and remediation.[4]

The Technology

Our product used an interactive speech technology system developed by Dragon Systems, Inc. The system included speech recognition software (which was integrated with our application software) and a printed circuit board with speech output as well as speech input capabilities. The high-quality speech output, achieved using custom-encoded LPC (linear predictive coding) speech synthesis, was produced with a Texas Instruments chip. The system ran on an Apple II+ or IIe computer with 64K and one disk drive.

The Dragon system could recognize up to thirty-two words at any given time. It was speaker dependent—meaning that each speaker had to "train" the system to recognize his or her pronunciation of the words in its vocabulary. And it recognized only isolated, not connected, speech—meaning that each word had to be preceded and followed by a brief silence.

The great advantage of the Dragon system was that, compared to other low-cost speech recognizers (and even some not so low-cost systems), it had great reliability. That is, it did not make many errors in recognition. Using a standard test of isolated word recognition, the system performed with 99.3 percent accuracy. Since the standard test was performed under optimum conditions, however, one object of our project was to determine whether the system was sufficiently reliable when used by young children in typical school settings to permit its use in early reading instruction.

The Early Reading Software

We developed and tested four software components:

- A training program, which both introduced the child to the recognizer's initial eighteen-word vocabulary (divided into two chunks of ten and eight words) and allowed the recognizer to collect the speech samples needed to run the other programs.[5] Clever graphics and music made the training program fun for children to use.
- A story program, entitled "Kody's Jungle Adventure," which helped the child to read the initial vocabulary in a meaningful and enjoyable context.
- A set of three games designed to help the child master the reading vocabulary.
- A program entitled "Your Words," which enabled a teacher or parent to enter additional vocabularies (of up to sixteen words each) which children could then train and use to play the games.

All components provided instructions and prompts using speech output; even the program

menus used speech as well as readily recognizable icons. Professional artist Jim Carson created the programs' visually appealing graphics, and original music was composed by a talented young musician, Paul Tegels.

The following sections describe the story program and the games in some more detail.

Kody's Jungle Adventure

The story program, which we developed with the guidance of reading specialist Carol Chomsky of the Harvard Graduate School of Education, was designed to give children practice in reading words in a meaningful context. The program featured a friendly koala named Kody. Speech output set the scene—Kody was lost in the jungle—and provided simple instructions: The child could help Kody find his way home by reading the words highlighted on the screen. A series of two-part episodes then presented Kody with a variety of obstacles (such as an elephant, a zebra, a tree, or a flower) and options for action (such as riding, tickling, climbing, or jumping). By reading the words highlighted on the screen, the child got the obstacles to appear and then got Kody to perform the actions.

For instance, the message "Here comes the elephant." might appear, with "elephant" highlighted. The program would say, "Here comes the." If the child said "elephant," the elephant would appear on the screen (see figure 1), accompanied by a musical theme.

Next, the message "Kody wants to lift it." might appear, with "lift" highlighted. The program would say, "Kody wants to." When the child said "lift," Kody would lift the elephant, and the music would play.

If the child said something other than the word called for, the program would suggest that she "Try again." If the child still could not read the word, the graphic of the obstacle or action would appear, and she would get another chance to read the word. The graphics therefore served as prompts as well as rewards for successful performance. If the third attempt failed, the program would read the entire sentence and then proceed.

At the end of each episode (obstacle plus action), a graphic showing Kody's progress through the jungle appeared at the top of the screen, and he was shown running toward home. After four or

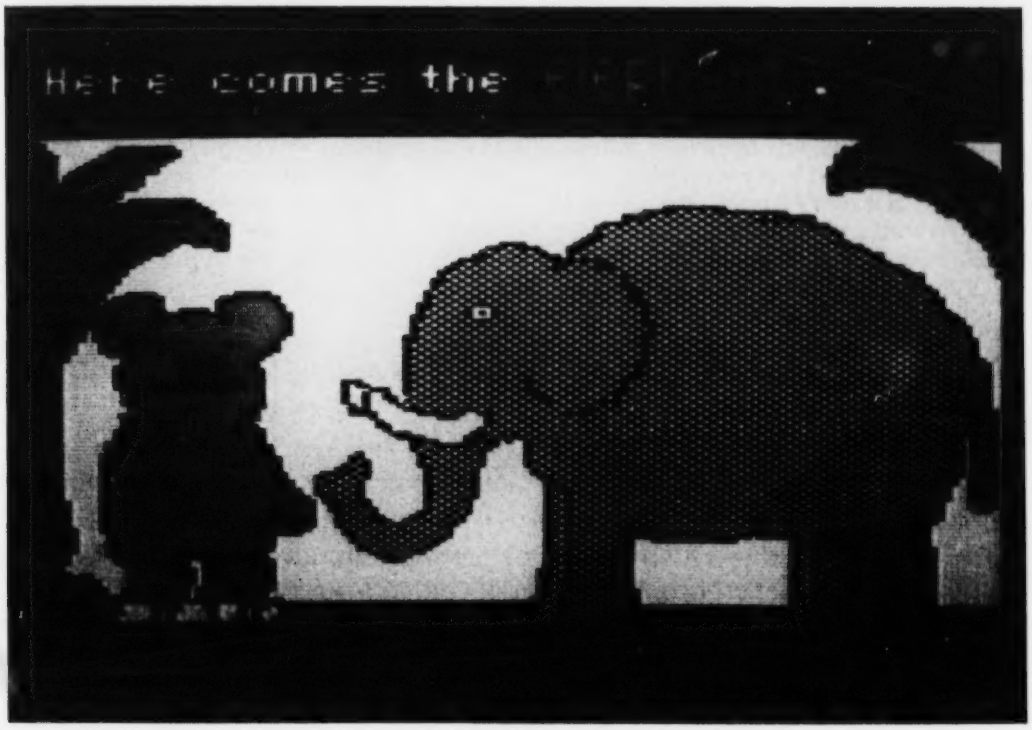

Figure 1.

BEST COPY AVAILABLE

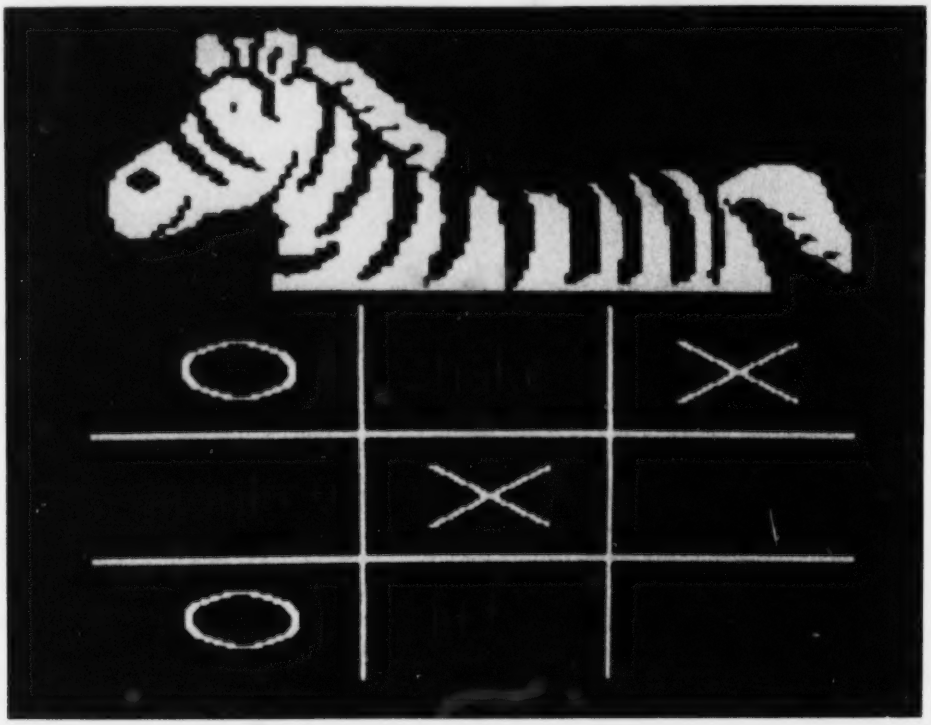

Figure 2.

five episodes (depending on the number of words trained), he would reach home. Each time the child read the story, it would consist of a different set of episodes; a total of twenty-eight combinations of obstacles and actions were available.

The story program was modeled on the way a parent or teacher may read a story to a child. The adult will often stop at various points and ask the child to read the next word. Of course, we had no intention of replacing such adult-child interaction; rather, we wanted to give children additional practice in recognizing words in meaningful contexts. We also expected the program to require some adult supervision, especially when first used. We felt that adults—or older children—should interact with youngsters learning to read, at the computer as well as away from it.

As for the story vocabulary, all the words were in the kindergarten to second grade range. Because our goal was to demonstrate the speech recognizer's capabilities, rather than to supplement any particular classroom reading program, we made no attempt to correlate the vocabulary with standard reading texts. Instead, we aimed for a range of difficulty and for words that would make the program interesting to children.

The Games

The software included three reading games:

Elephant Game (Rearrange the Words): Words appeared in a 3 x 3 grid, with one blank square. The child could move any word adjacent to the blank square into that square by reading it. The object was to move the word in the bottom lefthand corner to the top right—thereby letting the elephant out a door at the top of the grid. This game provided a lot of practice in reading the words, as most words had to be moved more than once.

Zebra Game (Tic Tac Toe): The child read the word in the box he wanted to mark with an X; the program played O (see figure 2). When the child won, the zebra danced.

Monkey Game (Concentration): Words appeared in a 3 x 3 grid. The child could reveal the colored rectangles "behind" a word by reading it, and then matched pairs of rectangles to reveal a piece of a picture of a monkey. When all the pairs were matched, the picture of the monkey was complete.

The games were independent of the story. Children could play the games with the story vocabulary or with any other set of words that they trained the program to recognize. (These additional sets could consist of anywhere from seven to sixteen words.) A number of additional vocabularies were provided, or teachers and parents could use their own word lists. Even children who had mastered the story vocabulary enjoyed practicing their reading skills by playing the games, and they could be further challenged by more advanced vocabularies.

Field Testing and Revising the System

We conducted three field tests of the system over a period of about six months with a total of forty-two kindergarten and first-grade children in two Boston-area public schools. In the first two field tests, which took place in Watertown, Massachusetts, children were brought in pairs to the school gymnasium, where we had set up the system. An adult introduced the program and helped the children as needed. An observer and video camera operator were also present. The third field test, in Newton, Massachusetts, took place in the school's computer laboratory, a setting more typical of actual school computer use. Students again worked in pairs with an adult, and the session was likewise observed and videotaped. Before each session, the adult administered a pretest to see if the children could read the words in the story vocabulary before encountering the program. Children were tested on these words again after the session. Participants in all three tests included students of low, average, and high ability as rated by their teachers.

The children who tried out the system in these field tests enjoyed using it, despite some technical glitches—especially in the first two trials.[6] Our field testing showed that the opportunity the program provided for the child to control what happened on the screen by reading words aloud was a powerful motivation to learn to read. The program proved especially successful with younger children (ages four to five) of average to high ability who were just beginning to recognize words, and with slightly older children (ages six to seven) who were having difficulty learning to read. Tables 1-3 (below and on page 6) show the mean numbers of words learned by children at each of the three trials.

Ironically, the technical problems with recognition, which were most severe in our first trial, seemed to account for the children learning to read more of the words in that trial. That is, because children had to repeat the correct words more times—both to train the system and to read the story and play the games—they had more practice with them, which probably explains why, according to the pretests and posttests we conducted, they learned more words.

After each trial, we worked to improve both the reading software and the recognition software. In revising the reading software, we focused particularly on the nature of the child-machine interaction—the dialogue that occurred between

Table 1. Number of words (out of ten) learned in Trial #1
Kindergarteners - 10th month of the school year

Teacher Rating of Ability	Mean Age Yrs-Mos	Mean Pretest Score	Mean Posttest Score	Mean Words Learned (Range)
Low (n=5)	5-7	0.0	3.6	3.5 (2-6)
Average (n=5)	6-0	6.4	6.4	6.4 (5-8)
High (n=7)	5-8	7.7	7.7	6.6 (4-8)

Table 2. Number of words (out of ten) learned in Trial #2
Kindergarteners - 3rd month of the school year

Teacher Rating of Ability	Mean Age Yrs-Mos	Mean Pretest Score	Mean Posttest Score	Mean Words Learned (Range)
Low (n=1)	5-2	0.0	0.0	0.0
Average (n=2)	4-2	0.0	3.0	3.0 (2-4)
High (n=3)	5-3	1.3	5.0	3.7 (2-8)
Unrated (n=1)	5-4	0.0	2.0	2.0

First Graders - 3rd month of the school year

Low (n=2)	6-6	0.5	7.5	7.0 (6-8)
Average (n=2)	6-8	2.0	8.5	6.5 (6-7)
High (n=4)	6-6	8.5	10.0	1.5 (0-4)

Table 3. Number of words (out of ten) learned in Trial #3
Kindergarteners - 4th month of the school year

Teacher Rating of Ability	Mean Age Yrs-Mos	Mean Pretest Score	Mean Posttest Score	Mean Words Learned (Range)
Low (n=3)	5-4	2.7	2.3	0.0
Average (n=2)	5-0	0.5	2.5	2.0 (1-3)
High (n=5)	5-3	3.4	7.4	4.0 (2-8)

the child and the program. Based on the kinds and frequency of prompts provided by the adult helping the children to use the program in the field tests, we added and, in some cases, removed prompts from the program. For example, we found it was useful for the recognizer to tell a child when his or her speech was too loud or too soft to be recognized. We then noted that we needed at least two different prompts for each case (for instance, "Louder, please," and "I can't hear you") to avoid monotonous repetition.

In adjusting the recognition software, we focused on reducing the frequency of errors, which were of three major kinds: rejection (rejecting a correct word), substitution (accepting an incorrect word), and insertion (responding to background noise, throat clearing, and other nonspeech sounds).

Because we couldn't eliminate recognition errors entirely, we also made efforts to adjust the relative frequency of different kinds of errors. For instance, we found that most children were quite willing to repeat a correct word—several times if necessary—if the recognizer didn't accept it for some reason. But when the recognizer accepted an incorrect word, they were either misled or confused. We therefore adjusted the software to make it more likely that the program would make rejection errors (rejecting correct words) than substitution errors (accepting incorrect words).

We also experimented with different kinds of microphones. Initially, we used an inexpensive, hand-held microphone. However, the children tended to play with this microphone, twisting the cord around their hands, lifting the microphone over their heads, and so on. They consequently did not hold it a uniform distance from their mouths, which impaired the effectiveness of the recognizer.

We then put this microphone in a stand, but the uniform distance problem wasn't solved, as the children bobbed their heads up and down, and back and forth. Finally, we settled on a headset microphone, which kept the distance from mouth to receiver constant.[7]

Overall, the recognizer's rate of accuracy—though significantly lower than the 99.3 percent achieved under optimal conditions—seemed acceptable. In the games, where the recognizer had a more difficult task, needing to distinguish among several acceptable words at any time, the mean number of errors per game ranged from 4 to 5.63. In the story, where the recognizer only had to identify one correct word at any time, the mean number of errors per story ranged from 1.75 to 4.60.[8]

In the second two trials, after we had made the adjustment described above, nearly all of the errors were rejection errors—that is, the child said the correct word and the recognizer rejected it. Most errors occurred with certain words, possibly because of quirks in the recognizer. For example, the long vowel sound in the word "feed" seemed to cause the recognizer particular difficulty. Also, in line with the conventional wisdom in the field that 80 percent of speech recognition errors occur among 20 percent of users, a small number of children tended to have large numbers of errors, thus skewing the mean error data.

Even at the end of our field tests, we were still experimenting with the microphones and adjusting the speech software and hardware. For example, we were fine-tuning the software-controlled gain setting, which affects sound amplification and therefore the recognizer's responses to background versus foreground noise. We were also able to use a separate hardware adjustment to raise the upper volume level that the recognizer would accept. We anticipated that ultimately we would be able to reduce the recognition error rate significantly.

Implications for Future Work

Probably the main lesson from our experience is that decisions related to human factors—everything from the kind of microphone used to the kinds of prompts and feedback the software provides—are critical in determining the likely success or failure of speech recognition with children. Moreover, the speech technology (both recognition and output) presents challenges beyond those that designers of educational software typically must deal with.

For example, in our use of speech input together with speech output, we were in effect attempting to simulate human dialogue. Some of the

kinds of difficulties this entails are illustrated by our efforts to design an effective training procedure. Since the child had to repeat each word several times during the training, the program somehow had to get the child to keep saying the word, but also to pause after each repetition (since the recognizer could handle only isolated words). In our initial design, the program prompted the child by saying, "Please say snake." (or whatever the word was) for each iteration. Each repetition was also signaled by a change in the accompanying graphic. We expected children to begin to anticipate the prompt and to use the graphic clue. However, they continued to wait for the prompt, which slowed the training process considerably.

In the second version, we therefore delayed the prompt (after the first instance of it), hoping that children would take the silence, along with the graphic, as a cue to proceed. But children continued to wait for the prompt, slowing the training down even further, rather than speeding it up.

Finally, we added an introductory segment to the training that modeled the repetition of words and explicitly asked children to attend to the moving word on the screen. This explicit approach proved effective and eliminated the need for repeated verbal prompts.

In addition to such challenges introduced by the attempt to simulate human dialogue, other issues arise from limitations in both speech recognition and speech output technologies. The need for some adult supervision—at least in children's initial encounters with the system and perhaps continuing—is a product of an inherent limitation of speech recognition systems that fall anywhere short of perfect performance. That is, unlike a person, the program cannot distinguish between recognizer error and child error. It cannot therefore encourage a child to repeat a correct word, in cases in which the recognizer rejects a correct response. Nor can it encourage a child to try another word when the error has been the child's rather than the recognizer's. The closest we came to being able to distinguish between human and recognizer error was in the story program, where we were able to provide different feedback depending on whether the recognizer considered a response to be close to the correct word or far from it.

While the main problem with speech recognition technology is insufficient discrimination, the main problem with speech output technology is inflexibility. We found that children varied considerably in their need for, and tolerance of, verbal prompts. At one extreme, some children needed not only the encouragement of the computer prompt to respond to the program but also the nodded or spoken approval of the adult who worked with the children in all of our trials. Other children, by contrast, were extremely impatient with prompts and spoke over them. But the program provided essentially the same quantity and kinds of prompts to all children alike. (This is, in fact, another example of the way the system inadequately replicated human dialogue.)

Such speech output problems may be amenable to technical solutions. Our system, for example, could not "listen" for speech while it was speaking, so a prompt, once begun, could not be interrupted. More advanced systems may be able to speak and listen simultaneously. We can also envision a system that would learn—on the basis of a user's typical speed of response—whether to prompt more or less. (We did add time-delayed repetitions of some of our speech prompts that occurred when there were long pauses before responses. When the adult present was instructed not to provide encouragement, these delayed prompts proved effective.)

As the preceding examples suggest, some of the problems we encountered may be inherent in the technology and therefore require instructional solutions external to the programs. Others likely have technological solutions. Perfecting both sorts of solutions requires extensive testing with users in real settings.

Despite the difficulties, however, the demonstration programs we produced showed that speech recognition technology has the potential to become a useful adjunct to early reading instruction. Students who used the system in our field tests all responded favorably to it. Even those who were accustomed to using computers at home or at school were fascinated by a computer that could both listen and talk. The programs' high level of interactivity kept stu-

dents engaged, as our videotapes of the testing sessions show, even in the early tests when the system did not always run smoothly. Though the small number of students tested and their limited use of the system preclude strong claims for student learning, the learning gains we did see at least suggest that such a system holds promise.

The positive results we achieved with a primitive speech recognition system and a computer with extremely limited memory and processing power underscore the potential of more advanced technology. Since the mid-1980s, when we did our work, speech recognition systems with much larger vocabularies, continuous speech capabilities, and much easier training requirements have been developed. But these systems are either quite costly or still in the laboratory. In fact, most commercially available speech recognizers today still have many of the same limitations as the system we used. Even DragonDictate, the costly voice-activated "typewriter" currently marketed by Dragon Systems to businesses, though it has a large vocabulary and the ability to learn a speaker's voice through use rather than a separate training procedure, cannot handle continuous speech. As the technology improves and as prices come down, however, our experiment points the way toward applications of speech recognition that go beyond drilling students on reading vocabulary.

Acknowledgements

The software this paper describes was a collaborative effort in which I was joined by Charles L. Thompson, then Director of the Center for Learning Technology at EDC; John H. McCloskey, whose programming genius made it all possible; and a number of people at Dragon Systems, who supplied the speech recognizer and served as technical resources on a variety of issues. They included Janet M. Baker, President of Dragon Systems; Mark Sidell, Principal Hardware Engineer; and Jed Roberts and David Pinto, Research Scientists. David Nelson, Director of Media Services at EDC, conducted the videotaping of the test sessions. The staff and students of the Phillips School in Watertown, Massachusetts, the Lincoln-Eliot School in Newton, Massachusetts, and the Plowshares Childcare Program in Newton, Massachusetts, participated in the field tests. The Watertown school system kindly loaned us videotaping equipment, and the Educational Collaborative for Greater Boston provided space for the testing. The research upon which this paper draws was conducted largely by Beth Wilson, with Charles L. Thompson and Philip P. Zodhiates, as part of EDC's subcontract to the Educational Technology Center at the Harvard Graduate School of Education. For their excellent suggestions for improving this paper, I would especially like to thank Leigh Peake and Cindy Char.

Notes

1. At the high end, for example, a voice-activated "typewriter," DragonDictate, with a basic vocabulary of 30,000 words, sells for $9,000 and requires an IBM PC AT compatible computer with 8 megabytes of memory, at an additional cost of $7,000 to $11,000 ("New 'Typewriter'" 1990; Rosenberg 1990). Even a far less sophisticated speech recognition system, the Voice Navigator from Articulate Systems, which runs on the Macintosh and recognizes 200 words at a time, sells for $1295.

2. A small number of speech recognizers marketed to schools, and an even smaller number of educational applications, are cited by Parham (1988) and Strickland et al. (1987).

3. The Chatterbox voice input card, which is available for both Apple II series and IBM PC compatible computers, sells for about $600; the software, which costs $245, also requires an Echo speech synthesizer ($129 for the Apple version) and a headphone/microphone ($75).

4. Much of the information in this paper is drawn from the report of a small research study conducted by the Educational Technology Center at the Harvard Graduate School of Education, which looked at some of the reliability and human interface issues raised by the recognizer (Thompson et al. 1985).

5. Although the system could actually handle up to thirty-two words at one time, we limited the vocabulary to eighteen words and then split it in two parts because of the need to have children train all the words used. In fact, we found that young children were willing to train ten words at one sitting, but not all eighteen. Beyond the initial eighteen-word vocabulary, however, additional vocabularies could be created by teachers or parents and trained and used by children. (See the discussion of additional vocabularies in the section entitled "The Games.")

6. In addition to problems with the speech recognition software and the microphones (described below), we also experienced some failures of the entire system due to static electricity.

7. However, at the time we stopped working on the product we were still seeking a headset microphone of sufficient quality and sufficiently low price that would fit well on child-sized heads.

8. Mean errors were calculated by adding the total number of recognition errors of all types by all children reading the story or playing a game during one of the trials and then dividing the total number of errors by the number of children who read a story or played a game. The figures therefore do not accurately reflect the number of errors for any one child reading the story or playing a game. Moreover, the total number of utterances, which varied widely (especially in the games) was not taken into account in calculating the mean errors. Selected data, for example, show that the total number of utterances per game could range from as low as 8 to as high as 61.

References

Blanchard, J. S., G. E. Mason, & D. Daniel. 1987. *Computer applications in reading*. Newark, DE: International Reading Association.

Evans, R. 1988 (November/December). The shock of recognition: Speech recognition is finally happening. *Language Technology* (10), 20-25.

New "typewriter" operated by voice. 1990 (March 20). *The Boston Globe*.

Parham, C. 1988. Computers that talk. *Classroom Computer Learning, 8* (6), 26-36, 63.

Rosenberg, R. 1990 (January 9). Crafting a new voice for words. *The Boston Globe*, p. 2.

Strickland, D. S., J. T. Feeley, & S. B. Wepner. 1987. *Using computers in the teaching of reading*. New York: Teachers College Press.

Thompson, C. L., B. Wilson, P. P. Zodhiates, & I. Kantrov. 1985. Speech recognition technology: An application to beginning reading instruction. Technical Report No. 85-7, Educational Technology Center, Harvard Graduate School of Education, Cambridge,MA.